

### Linux, Energy, and Networks:

"Saving Large Amounts of Energy With Network Connectivity Proxying"

#### **Bruce Nordman**

**Lawrence Berkeley National Laboratory** April 9, 2009

BNordman@LBL.gov — efficientnetworks.LBL.gov

RENCE BERKELEY NATIONAL LABORATORY

ENERGY FFICIENT DIGITAL

NETWORKS

Linux and Energy Efficiency\*



#### Five basic dimensions of topic

- Efficiency of computing a lot
- Efficiency of computing a little
- Efficiency of doing no computing
- Effectiveness of communicating with user
- Efficiency imposed on other devices (via network)

\*Efficiency not in traditional physics terms



Slide 2 of 39

#### **Need to Think Broadly about Networks**



















Slide 3 of 39

LAWRENCE BERKELEY NATIONAL LABORATORY

#### How much energy does The Internet use?

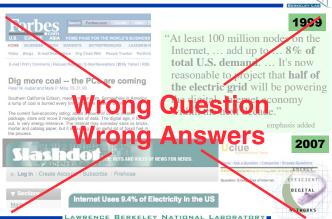




WRENCE BERKELEY NATIONAL LABORATORY

#### How much energy does The Internet use?





#### **Network Structure and Energy**

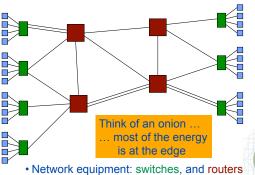


ENERGY FFICIENT

DIGITAL

NETWORKS

• Edge devices: PCs, servers - Displays, storage, phones, ...



LAWRENCE BERKELEY NATIONAL LABORATORY Slide 6 of 39

#### **Networks and Energy**



Product

Network Int.

Network

Product

Link

## Network equipment .... Routers, switches, modems, wireless APs, ... ... vs networked equipment

PCs, printers, set-top boxes, ...

#### How networks drive energy use

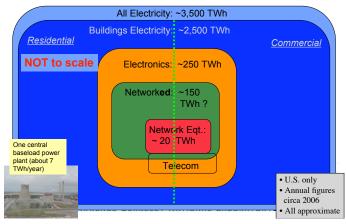
- Direct
  - -Network interfaces (NICs)
  - -Network products
- Induced in Networked products
  - -Increased power levels
  - Increased time in higher power modes (to maintain network presence)

Slide 7 of 39

LAWRENCE BERKELEY NATIONAL LABORATORY

#### **Network electricity use in context**





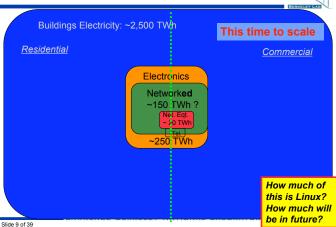
#### Network electricity use in context, cont.



ENERGY

FFICIENT

DIGITAL



#### How to think about energy quantities



Our needs only require approximations

1 year = 8,760 hours ~ 10,000 hours

1 kWh costs \$0.09 ~ \$0.10

1 W for 1 year ~ \$1

1 TWh = 1 billion kWh ~ \$100 million

U.S. annual consumption ~ 3,500 TWh

... buildings portion ~ 2,500 TWh



Slide 10 of 39

LAWRENCE BERKELEY NATIONAL LABORATORY

## Things we know: Energy consumption is at edge



- Network equipment < 10% of all electronics
- Most electronics already networked
- More electronic and non-electronic devices getting networked
- Network induced consumption > all direct
- · Network equipment energy will grow ...
- ... but other electronics will grow faster

#### Things we know: Utilization is low



Data networks are lightly utilized, and will stay that way,
A. M. Odlyzko, Review of Network Economics, 2003

Network

Utilization

NetworkUtilizatioAT&T switched voice33%Internet backbones15%Private line networks3~5%LANs1%

To brow buddeth Uktuba Profile

Low utilization is norm in life — e.g. cars

- Average U.S. car ~12,000 miles/year = 1.5 miles/hour
- If capacity is 75 mph, this is 2% utilization

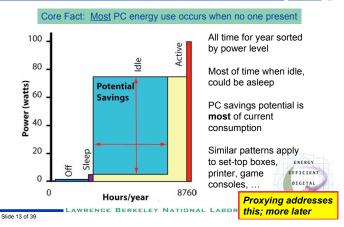
PC and server utilization also low



ENERGY
EFFICIENT
DIGITAL
NETWORKS

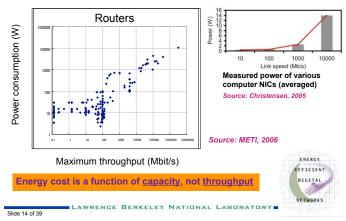
## Things we know: Edge device energy is mostly idle





## Things we know: Speed costs energy / power





## Things we know: IP will go everywhere



- IT equipment IP already universal
- IP for phone calls (VOIP)
- IP for TV (IPTV)
- IP for consumer electronics generally
- IP for buildings (lighting, climate)
- IP for .....

How much of this will be Linux?

ENERGY

Slide 15 of 39

#### **Efficiency Approaches**



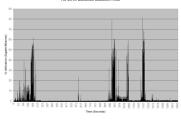
#### **Product** Network Interface Protocol / **Focus Product Focus Application Focus Focus** Examples: Energy FFICIENT Efficient **Proxying** CE DIGITA NETWORKS AWRENCE BERKELEY NATIONAL LABORATORY Slide 16 of 39

#### Adaptive Link Rate (ALR)



FFICIENT

DIGITAL



16 14 10 100 1000 10000 10000 10000 Link speed (Mb/s)

#### Observations

- Most of time, full link capacity not needed
- Notebooks already dropped link rate in sleep

Proposal (LBNL & USF)

 Enable changing link rate quickly in response to traffic levels (ms not s)

 ENERGY.

 ENERGY.

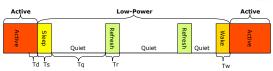
#### (ALR now) Energy Efficient Ethernet



Energy

**TEfficient** 

Ethernet



- IEEE 802.3az created to standardize EEE
- Standards process began with ALR; eventually settled on alternate method "Low Power Idle"
  - Stop transmitting between packetsSwitch now takes *micro*seconds
- · Standards process needs about 1 more year
  - Goal to get EEE technology into ALL Ethernet network hardware globally over next few years









#### Our CE Future?





While some integrators are skeptical about the prewired preprogrammed NHS rack from Sony, others embrace the solution for its simplicity.

- Network / Data connectivity a Mess
- · Number of CE devices is LARGE
- · For energy use, digital networking could easily:
  - cause large increases, or
  - enable significant reductions
- · We cannot rely on manual power control

#### ENERGY FFICIENT DIGITAL

#### Consumer Electronics – What to do



- · Move to 3-state power model
- · Address link power consumption
- Provide for persistent network presence
- Expose power state to network
- · Standardize some user interface elements
  - Displays
- · Create a model for standard behaviors / expectations for CE devices

ENERGY

Many of these devices will run Linux. Any implications for OS or related activities?

Slide 21 of 39

LAWRENCE BERKELEY NATIONAL LA

#### **User Interfaces**



- Standard Interface elements common throughout daily life
- Key to safety, ease of use, efficiency
- Many use graphics, color, location, etc. to improve functionality and reduce languagedependence
- Commonality limited to comprehension needs
- Can deviate from standards when there is a good reason



AWRENCE BERKELEY NATIONAL LABORATORY



Slide 22 of 39

#### **User Interface Standards**

**Key Elements** 

• Terms

· Colors

Symbols

Metaphors



- · Consistent across:
- Manufacturers
- Products
- Countries
- Simple
- Accessible
- Portable









LAWRENCE BERKELEY NATIONAL LABORATORY

#### Non-Interoperability w/ devices or w/ people



· Failure to accomplish interoperability:



- Is annoying

- Costs product manufacturers
  - Design-----
  - Manufacture % Sales 1)
- Wastes energy 1
  - · Difficult or impossible to match wanted service to delivered
- Impedes addressing climate change



ENERGY FFICIENT DIGITAL

Slide 24 of 39

#### **User Interfaces**



#### People:

- ... are best understood as nodes on the network
  - Even more than portable electronics, they move
- · ... are often absent from design, presentation of networks
- · ... need standard interfaces, just like devices do
  - Nature of interface different, but principle same

Past LBNL work: "Power Control User Interface Standard", IEEE 1621 - terms, symbols, colors, metaphor



Linux community should adopt IEEE 1621

ENERGY

EFFICIENT

Slide 25 of 39

LAWRENCE BERKELEY NATIONAL LABORAT

#### "Network Connectivity Proxying"

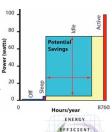


A low-power entity that maintains "full" network connectivity for a sleeping high-power device

- Addresses energy use by devices persistently network-connected, but often doing little or nothing
- Key goal: hide host's sleep state from rest of network
- · Need standard definition of proxy behavior
- · Need cooperation of operating system

Key collaborator: Ken Christensen, University of South Florida

LAWRENCE BERKELEY NATIONAL LABORA



Linux community should engage proxying topic

### Networks and Energy



Network equipment ....
Routers, switches, modems, wireless APs, ...
... vs networked equipment
PCs, printers, set-top boxes, ...

How networks drive energy use

- Direct
  - Network interfaces (NICs)
- Network products

Induced in Networked products

- Increased power levels
- Increase time in higher power modes
- (to maintain network presence)

Network Product

ENERGY

EFFICIENT
DIGITAL

NETWORKS

Product

Network Int.

Slide 27 of 39

LAWRENCE BERKELEY NATIONAL LABORATORY

#### **Proxying: Origins**



#### 3.4 PC Power Management with Networks

Networks pose special challenges for power management. Depending on the systems (hardware and software), the network can partially or entirely defeat power management, or may require extra configuration changes for it to function.

LAWRENCE BERKELEY NATIONAL LABORATORY

INTERNATIONAL JOURNAL OF NETWORK MANAGEMENT Int. J. Network Mgmt., 8, 120–130 (1998)

Enabling Power Management for Network-attached Computers

Power management is an emerging area of interest for network management. This article reviews current developments and describes methods for enabling power management in network-attached computers. © 1999 John

By Kenneth J. Christensen\* and Franklin 'Bo' Gulledge

Slide 28 of 39

Slide 30 of 39

Slide 26 of 39

LBNL Report: 1997

USF paper: 1998

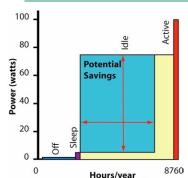
# ENERGY EFFICIENT DIGITAL NETRORKS

#### PC energy is mostly idle



#### Core Fact: Most PC energy use occurs when no one present

AWRENCE BERKELEY NATIONAL LABORATORY



All time for year sorted by power level

Most of time when idle, could be asleep

PC savings potential is **most** of current consumption

Similar patterns apply to set-top boxes, printer, game consoles, ...

DIGITAL

## Proxy does: ARP, DHCP, TCP, ICMP, SNMP, SIP, ....

·**,** ····· , ···· , ···· , ···· , ···· , ···· , ···· , ···· , ····

**Proxying: Operation** 

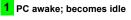
Proxy

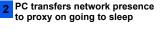
BERKELEY LAD

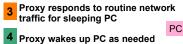
LAN or

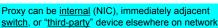
Internet

#### Proxy operation











LAWRENCE BERKELEY NATIONAL LABORATORY

Slide 29 of 39

#### **Proxying: Relevant Protocols**



- · What is network presence?
  - -Host-level reachability
    - ARP. IGMP
  - -Application-level reachability
    - · TCP SYN, SIP invitations
  - -Addressability
    - DHCP
  - -Manageability
    - · ICMP, SNMP
  - -l iveliness
    - · TCP connections, application heartbeats



LAWRENCE BERKELEY NATIONAL LABORATORY Slide 31 of 39

#### **General Goals**



- Enable large majority of PC users to use sleep without breaking their own or IT admin applications
  - At least 80%. > 90% better. > 95% or > 98% even better.
- Enable both current and emerging common applications
- Enable standard to directly (or easily adapted) for use in printers, set-top boxes, game consoles, etc.
- Describe behavior of "green applications" that don't break proxying
  - · Create de facto guide for new applications

#### **General Goals, restated**



#### PC (or other edge device):

- Is always available
- · Doesn't wake up if doesn't need to
- Does wake up when does need to
- Provides good user experience
- · Provides consistent user experience
- Hides sleep status from rest of network
  - · Except when explicitly tells



Slide 33 of 39

LAWRENCE BERKELEY NATIONAL LABORATORY

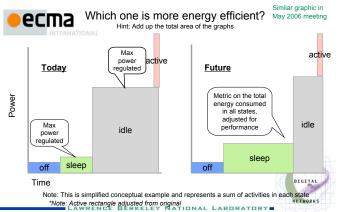
#### Ecma Context - TC38-TG2 (May, 2007)



ENERGY

EFFICIENT

DIGITAL NETWORKS



Slide 34 of 39

Slide 32 of 39

#### **Energy Star context**



Computer Performance and Energy Assessment Tool Stakeholder Meeting, June 20, 2007, Washington, D.C.



#### Background

- ckground
  Most energy used by desktop PCs in U.S. when no one is present
  Enabling power management could save > 50% of desktop PC energy use
  Network connectivity the major impediment to enabling sleep moving forward
  Topic dates back to beginning of Energy Star PC process in 2004
  Intent is to enable sleep without requiring any changes to existing protocols and
  applications used on great majority of PCs
  Wake On LAN inadequate for general solution for many reasons

Drive proxying functionality into all networked electronic products that have significant

EPA Announcement of V4.0 Process, September 2004

I) Fix the "network problem" with power management

In future, Linux community should become more engaged with Energy Star

#### Energy Star V5.0 Computer Spec (eff. June, 2009)



#### Definition (emphasis added)

Full Network Connectivity: The ability of the computer to maintain network presence while in sleep and intelligently wake when further processing is required. Maintaining network presence may include obtaining and/or defending an assigned interface or network address, responding to requests from other nodes on the network, or sending periodic network presence messages to the network all while in the sleep state. In this fashion, presence of the computer, its network services and applications, is maintained even though the computer is in sleep.

#### Requirement: None

#### Incentive

Reduced idle time in TEC calculation



#### **Proxying: Process**





Slide 37 of 39

LAWRENCE BERKELEY NATIONAL LABORATORY

#### Thank you!



efficientnetworks.LBL.gov
Bruce Nordman

Bruce Nordman Lawrence Berkeley National Laboratory BNordman@LBL.gov 510-486-7089

(or google)





Slide 39 of 39

LAWRENCE BERKELEY NATIONAL L





#### **Key Points**



- Standard intended for PCs and any "PC-like" device
  - -Desire for persistent network connectivity
  - -Non-trivial difference between idle/sleep power
  - -PCs, printers, set-top boxes, game consoles, ...
- · Establishes a floor of functionality, not a ceiling
- · Not designed with servers in mind
- Avoid any content that limits location of proxy between Internal (NIC), and External (closest switch or router)
  - -NOT get distracted by "third-party" proxy location

Need contacts in Linux community to assure this gets implemented in timely fashion

Slide 38 of 39